



10 RELATED APPLICATIONS

15 PROCESSING ENVIRONMENT WITH SYNCHRONIZATION METHODS EMPLOYING A
UNIFICATION DATABASE; serial no. 60/094,972, filed July 31, 1998, and entitled SYSTEM
AND METHODS FOR SYNCHRONIZING TWO OR MORE DATASETS; and serial no.
60/094,824, filed July 31, 1998, and entitled DATA PROCESS ENVIRONMENT WITH
METHODS PROVIDING CONTEMPORANEOUS SYNCHRONIZATION OF TWO OR

25 09/136,215, filed August 18, 1998, now U.S. Patent 6,295,541, and entitled SYSTEM AND METHODS FOR SYNCHRONIZING TWO OR MORE DATASETS. The present application is also related to the following commonly owned U.S. patent applications, the disclosures of which are hereby incorporated by reference in their entirety, including any appendices or attachments thereof, for all purposes: serial no. 08/609,983, filed February 29, 1996, now U.S.

Page 1 of 25

[0006] As the use of information appliances is ever growing, often users will have their data in more than one device, or in more than one desktop application. Consider, for instance, a user who has his or her appointments on a desktop PC (personal computer) but also has a battery-powered, hand-held device for use in the field. What the user really wants is for the information of each device to remain synchronized with all other devices in a convenient, transparent manner. Still further, the desktop PC is typically connected to a server computer, which stores information for the user. The user would of course like the information on the server computer to participate in the synchronization, so that the server also remains synchronized.

[0007] A particular problem exists as to how one integrates disparate information -- such as calendaring, scheduling, and contact information -- among multiple devices, especially three or more devices. For example, a user might have a PalmPilot ("Pilot") device, a REX™ device, and a desktop application (e.g., Starfish Sidekick running on a desktop computer). Currently, in order to have all three synchronized, the user must follow a multi-step process. For instance, the user might first synchronize data from the REX™ device to the desktop application, followed by synchronizing data from the desktop application to the Pilot device. The user is not yet done, however. The user must synchronize the Pilot back to the REX™ device, to complete the loop. Description of the design and operation of the REX™ device itself (available as Model REX-3, from Franklin Electronic Publishers of Burlington, NJ) is provided in commonly-owned U.S. patent application serial no. 08/905,463, filed August 4, 1997, and entitled, USER INTERFACE METHODOLOGY FOR MICROPROCESSOR DEVICE HAVING LIMITED USER INPUT, the disclosure of which is hereby incorporated by reference.

[0008] Expectantly, the above point-to-point approach is disadvantageous. First, the approach requires user participation in multiple steps. This is not only time consuming but also error prone. Further, the user is required to purchase at least two products. Existing solutions today are tailored around a device-to-desktop PIM (Personal Information Manager) synchronization, with no product capable of supporting concurrent synchronization of three or more devices. Thus for a user having three or more devices, he or she must purchase two or more separate synchronization products. In essence, existing products to date only provide peer-to-peer synchronization between two points, such as between point A and point B. There is no product providing synchronization from, say, point A to point B to point C, all at the same time.

Instead, the user is required to perform the synchronization manually by synchronizing point A to point B, followed by synchronizing point B to point C, then followed by point C back to point A, for completing the loop.

[0009] As a related disadvantage, existing systems adopt what is, in essence, an approach having a "hard-coded" link for performing synchronization for a given type of data. Suppose, for example, that a user desires to update his or her synchronization system for now accommodating the synchronization of e-mail data (e.g., Microsoft® Outlook e-mail). With existing synchronization products, the user cannot simply plug in a new driver or module for supporting this new data type. To the point, existing products today do not provide a generic framework into which data type-specific modules may plug into. As a result, these products are inflexible. In the event that the user encounters a new type of data for which synchronization is desired, he or she is required to update all or substantially all of the synchronization product. The user cannot simply plug in a driver or module for supporting synchronization of the new data type. All told, existing synchronization products today assume that users will only perform point-to-point (i.e., two device) synchronization, such as between a hand-held device and a desktop application running on a PC.

[0010] This assumption is far removed from reality, however. Users are more likely today to have data among multiple devices, such as among a desktop computer, a server computer (e.g., company network at the user's place of employment), and two or more portable devices (e.g., a laptop computer and a hand-held device). Given the substantial effort required to manually keep three or more devices synchronized, the benefits of synchronization largely remain unrealized for most computer and information application users today.

[0011] What is needed is a system providing methods which allows a user of information processing devices to synchronize user information, such as user-supplied contact lists, from one device to any number of other devices, including three or more devices concurrently. The present invention fulfills this and other needs.

SUMMARY OF THE INVENTION

[0012] The present invention introduces the notion of a reference database: the Grand Unification Database or GUD. By storing the data that is actually being synchronized (i.e.,

storing the actual physical body of a memo, for instance) inside an extra database (or by specially-designated one of the client data sets) under control of a central or core synchronization engine, rather than transferring such data on a point-to-point basis, the system of the present invention provides a repository of information that is available at all times and does not require that any other synchronization client (e.g., PIM client or hand-held device) be connected.

Suppose, for instance, that a user has two synchronization clients: a first data set residing on a desktop computer and a second data set residing on a hand-held device. The GUD introduces a third data set, a middleware database. This third data set provides a super-set of the other two client data sets. Therefore, if the user now includes a third client, such as a server computer storing user information, the synchronization system of the present invention has all the information necessary for synchronizing the new client, regardless of whether any of the other clients are currently available. The system can, therefore, correctly propagate information to any appropriate client without having to "go back" to (i.e., connect to) the original client from which that data originated.

[0013] Internally, the system of the present invention employs "type plug-in" modules, each one for supporting a particular data type. Since the core synchronization engine treats data generically as "blob" objects, type-specific support is provided by the corresponding plug-in module. Each plug-in module is a type-specific module having an embedded record API (application programming interface) that each synchronization client may link to, for providing type-specific interpretation of blob data. For instance, the system may include one type-specific record API for contact information, another for calendar information, and yet another for memo information. In this manner, each client may employ a type-specific API for correctly interpreting and processing particular blob data. The engine, on the other hand, is concerned with correct propagation of data, not interpretation of that data. It therefore treats the data itself generically. In this fashion, the present invention provides a generic framework supporting concurrent synchronization of an arbitrary number of synchronization clients or devices.

[0014] Also internally, the synchronization system of the present invention employs an "action queue," for optimizing the actual synchronization work performed. In contrast to conventional point-to-point (i.e., binary) synchronization systems, the synchronization system of the present invention does not immediately transmit updates or changes as soon as they are

detected. Instead, the system determines or tabulates changes, net of all clients, before undertaking the actual work (e.g., record insertion) of synchronizing a particular client. In particular, all actions or tasks which are to be performed for a client by the system during synchronization are queued in the outbound action queue. This allows the system to apply synchronization logic or intelligence to the queue for further improving system performance, such as eliminating any activities which are redundant or moot. For example, if the system receives a request from two different clients to update a given record (i.e., conflict), the system, applying internal synchronization logic, can eliminate propagating the first update, as it is rendered moot by the second update. In this manner, the system can apply a first-level resolution of requests that are conflicting (or complimentary) and, as a result, eliminate those synchronization activities which are redundant or moot.

[0015] An exemplary method for synchronizing multiple data sets includes first establishing a data repository for facilitating synchronization of user information maintained among multiple data sets, the data repository storing user information from the data sets. At least one mapping is stored which specifies how user information may be transformed for storage at a given data set. Upon receiving a request for synchronizing at least one data set, the system may, based on user information stored at the data set(s) and based on the mapping, propagate to the data repository from each data set(s) any changes made to the user information, to the extent that such changes can be reconciled with user information already present at the data repository. Further, based on user information stored at said data repository and based on the mapping, the system may propagate to each data set(s) any changes to the user information which have been propagated to the data repository, to the extent that such changes are not present at the data set.

BRIEF DESCRIPTION OF THE DRAWINGS

[0016] Fig. 1A is a block diagram of a computer system in which the present invention may be embodied.

[0017] Fig. 1B is a block diagram of a software system of the present invention for controlling operation of the system of Fig. 1A.

[0018] Fig. 2 is a block diagram of the synchronization system of the present invention.

[0019] Fig. 3 is a block diagram of a GUD of the present invention.

[0020] Figs. 4A-C are flow charts of the operation of the synchronization system of the present invention.

DETAILED DESCRIPTION OF A PREFERRED EMBODIMENT

5 [0021] The following description will focus on the presently-preferred embodiment of the present invention, which is operative in an environment typically including desktop computers, server computers, and portable computing devices, occasionally or permanently connected to one another, where synchronization support is desired. The present invention, however, is not limited to any particular environment or device. Instead, those skilled in the art will find that the
10 present invention may be advantageously applied to any environment or application where contemporaneous synchronization among an arbitrary number of devices (i.e., "synchronization clients"), especially three or more devices, is desirable. The description of the exemplary embodiments which follows is, therefore, for the purpose of illustration and not limitation.

15 System hardware and software

[0022] The present invention may be embodied on an information processing system such as the system 100 of Fig. 1A, which comprises a central processor 101, a main memory 102, an input/output (I/O) controller 103, a keyboard 104, a pointing device 105 (e.g., mouse, pen device, or the like), a screen or display device 106, a mass storage 107 (e.g., hard disk, removable
20 floppy disk, optical disk, magneto-optical disk, flash memory, or the like), one or more optional output device(s) 108, and an interface 109. Although not shown separately, a real-time system clock is included with the system 100, in a conventional manner. The various components of the system 100 communicate through a system bus 110 or similar architecture. In addition, the system 100 may communicate with other devices through the interface or communication port
25 109, which may be an RS-232 serial port or the like. Devices which will be commonly connected to the interface 109 include a network 151 (e.g., LANs or the Internet), a laptop 152, a handheld organizer 154 (e.g., the REX™ organizer, available from Franklin Electronic Publishers of Burlington, NJ), a modem 153, and the like.

[0023] In operation, program logic (implementing the methodology described below) is
30 loaded from the storage device or mass storage 107 into the main memory 102, for execution by

the processor 101. During operation of the program (logic), the user enters commands through the keyboard 104 and/or pointing device 105 which is typically a mouse, a track ball, or the like. The computer system displays text and/or graphic images and other data on the display device 106, such as a cathode-ray tube or an LCD display. A hard copy of the displayed information, or other information within the system 100, may be obtained from the output device 108 (e.g., a printer). In a preferred embodiment, the computer system 100 includes an IBM PC-compatible personal computer (available from a variety of vendors, including IBM of Armonk, New York) running Windows 9x or Windows NT (available from Microsoft Corporation of Redmond, Washington). In a specific embodiment, the system 100 is an Internet or intranet or other type of network server and receives input from and sends output to a remote user via the interface 109 according to standard techniques and protocols.

[0024] Illustrated in Fig. 1B, a computer software system 120 is provided for directing the operation of the computer system 100. Software system 120, which is stored in system memory 102 and on storage (e.g., disk memory) 107, includes a kernel or operating system (OS) 140 and a windows shell 150. One or more application programs, such as client application software or "programs" 145 may be "loaded" (i.e., transferred from storage 107 into memory 102) for execution by the system 100.

[0025] System 120 includes a user interface (UI) 160, preferably a Graphical User Interface (GUI), for receiving user commands and data and for producing output to the user.

These inputs, in turn, may be acted upon by the system 100 in accordance with instructions from operating system module 140, windows module 150, and/or client application module(s) 145. The UI 160 also serves to display the user prompts and results of operation from the OS 140, windows 150, and application(s) 145, whereupon the user may supply additional inputs or terminate the session. In the preferred embodiment, OS 140 and windows 150 together comprise Microsoft Windows software (e.g., Windows 9x or Windows NT). Although shown conceptually as a separate module, the UI is typically provided by interaction of the application modules with the windows shell and the OS 140.

[0026] Of particular interest herein is a synchronization system or "Synchronizer" 200 of the present invention, which implements methodology for contemporaneous synchronization of an arbitrary number of devices or "clients." Before describing the detailed construction and

operation of the Synchronizer 200, it is helpful to first briefly review the basic application of synchronization to everyday computing tasks.

Brief overview of synchronization

A. Introduction

[0027] Many software applications, such as personal productivity applications as Starfish Sidekick® and Lotus® Organizer, have sets of data or “data sets” (e.g., address books and calendars). Consider for instance a user scenario where an account executive needs to coordinate contacts and events with other employees of the XYZ corporation. When traveling, this executive carries a laptop PC with Starfish Sidekick® installed. At home, she and her husband use Lotus® Organizer to plan their family's activities. When on family outings, the account executive carries her PalmPilot™ hand-held organizer. As the foregoing illustrates, a user often needs a means for synchronizing selected information from the data sets his or her applications rely upon. The account executive would not want to schedule a business meeting at the same time as a family event, for example.

[0028] Conventionally, the process of synchronizing or reconciling data sets has been a binary process -- that is, two logical data sets are synchronized at a time. Any arbitrary synchronization topology will be supported. Here, the system guarantees synchronization stability and the avoidance of undesirable side effects (cascading updates, record duplication, or the like). Data sets do not need to be directly connected but, instead, can be connected via a “store-and-forward” transport, such as electronic mail.

B. Synchronization design

1. Synchronization type

[0029] Data set synchronization may, for convenience of description, be divided into two types: content-oriented and record-oriented. Content-oriented synchronization correlates data set records based on the values of user-modifiable fields. Value correlation requires semantic (or at least advanced syntactic) processing that the human brain is very good at and computers are not. For example, a record in one data set with a name field valued “Johann S. Bach” and a record in a second data set with a name field valued “J. S. Bach” could possibly refer to the same real-world

person. A human being might arrive at this conclusion by correlating associated data (addresses) or drawing upon external information (e.g., Bach is an unusual name in the U.S.). Creating program logic or code with the ability to make these type of decisions is computationally very expensive.

5 **[0030]** Record-oriented synchronization correlates data set records by assuming that each record can be uniquely identified throughout its lifetime. This unique identifier is usually implemented as a non-modifiable, hidden field containing a "Record ID". Record-oriented synchronization algorithms usually require maintaining a mapping from one set of record IDs to another. In a preferred embodiment, the system employs record-oriented synchronization.

10 **[0031]** Record-oriented synchronization is conceptually simple and may be summarized as follows. In the rules below, A and B refer to two data sets which have a synchronization relationship. The rules are assumed to be symmetrical.

- 15 1. A and B must track similar types of data (e.g., if A is an address book, then B must be an address book).
2. A record entered in A, will create a record in B.
3. A record modified in A, will modify the corresponding record in B.
4. If record A1 has been modified in A and the corresponding record B1 has been modified in B, the record with the latest timestamp takes precedence.

20

The rules presented above reduce the occurrence of undesirable side effects with a network of synchronized data sets.

2. Timestamps

25 **[0032]** The actual synchronization logic in synchronization systems often needs to make processing decisions based on comparing the time at which past events occurred. For example, it is necessary to know if a record was modified before or after the last synchronization transaction. This requires recording the time of various events. A "timestamp" value may be employed to this purpose. Typically, data sets involved in synchronization support timestamps, or can be supplied with suitable timestamps, in a conventional manner. In conjunction with the usage of timestamps

30 to compare the relative timing of record creation or modification, the clocks on the respective devices may themselves be synchronized.

3. Record Transformations

[0033] During synchronization, a synchronization system will typically transform records from one application-usage-schema set to another application-usage-schema set, such as transforming from a Starfish Sidekick® card file for business contacts to a corresponding PalmPilot™ data set. Typically, there is a one-to-one relationship between records in these two data sets, that is, between the source and target data sets. If this is not the case, however, the component of the system that interacts with a non-conforming data set may include logic to handle this non-conformance.

[0034] The record transformations themselves are a combination of field mappings and conversions from a source record to a target record. Exemplary types of field mappings include, for instance, the following.

- | | |
|-----------------------|--|
| 1. <i>Null</i> | Source field has no equivalent field in the target data set and is ignored during synchronization. |
| 2. <i>One-to-One</i> | Map exactly one field in the target to one field in the source. |
| 3. <i>One-to-Many</i> | Map one field in the target to many fields in the source, such as parse a single address line to fields for number, direction, street, suite/apartment, or the like. |
| 4. <i>Many-to-One</i> | Map several fields in the target to one field in the source, such as reverse the address line mapping above. |

Similarly, exemplary field conversions may be defined as follows.

- | | |
|---------------------------|--|
| 1. <i>Size</i> | Source field may be larger or smaller in size than the target field. |
| 2. <i>Type</i> | Data types may be different, such as float/integer, character vs. numeric dates, or the like. |
| 3. <i>Discrete Values</i> | A field's values may be limited to a known set. These sets may be different from target to source and may be user defined. |

It is often the case that there are significant differences in the number, size, type and usage of fields between two data sets in a synchronization relationship. The specification of transformations is typically user-configurable, with the underlying system providing defaults.

[0035] With an understanding of the basic process of synchronizing information or computing devices, the reader may now better appreciate the teachings of the present invention for providing improved methodology for contemporaneous synchronization of an arbitrary

number of devices (i.e., synchronization clients). The following description focuses on specific modifications to a synchronization system for implementing the improved synchronization methodology.

5 **Synchronization system providing contemporaneous synchronization of two or more clients**

A. General design considerations

[0036] The present invention introduces the notion of a "Grand Unification Database" (GUD) -- a central repository or reference database for user data. By storing the data that is actually being synchronized (i.e., storing the actual physical body of a memo, for instance) inside
10 an extra database (or by specially-designated one of the client data sets) under control of a central or core synchronization engine, rather than transferring such data on a point-to-point basis, the system of the present invention provides a repository of information that is available at all times and does not require that any other synchronization client (e.g., PIM client or hand-held device) be connected. Suppose, for instance, that a user has two synchronization clients: a first data set
15 residing on a desktop computer and a second data set residing on a hand-held device. The GUD introduces a third data set, a middleware database. This third data set provides a super-set of the other two client data sets. Therefore, if the user now includes a third client, such as a server computer storing user information (or other information which the user desires synchronization to), the synchronization system of the present invention has all the information necessary for
20 synchronizing the new client, regardless of whether any of the other clients are currently available. The system can, therefore, correctly propagate information to any appropriate client without having to "go back" to (i.e., connect to) the original client from which that data originated.

[0037] Internally, the system of the present invention employs a driver-based architecture
25 providing type-specific "plug-in" modules, each one for supporting a particular data type. Since the core synchronization engine treats data generically as "blob" objects, type-specific support is provided by the corresponding plug-in module. Each plug-in module is a type-specific module having an embedded record API (application programming interface) that each synchronization client may link to, for providing type-specific interpretation of blob data. For instance, the
30 system may include one type-specific record API for contact information, another for calendar

information, and yet another for memo information. In this manner, each client may employ a type-specific API for correctly interpreting and processing particular blob data. The engine, on the other hand, is concerned with correct propagation of data, not interpretation of that data. It therefore treats the data itself generically. In this fashion, the present invention provides a generic framework supporting concurrent synchronization of an arbitrary number of synchronization clients or devices.

[0038] Also internally, the synchronization system of the present invention employs an "action queue," for optimizing the actual synchronization work performed. In contrast to conventional point-to-point (i.e., binary) synchronization systems, the synchronization system of the present invention does not immediately transmit updates or changes as soon as they are detected. Instead, the system determines or tabulates changes, net of all clients, before undertaking the actual work (e.g., record insertion) of synchronizing a particular client. In particular, all actions or tasks which are to be performed for a client by the system during synchronization are queued in the outbound action queue. This allows the system to apply synchronization logic or intelligence to the queue for further improving system performance, such as eliminating any activities which are redundant or moot. For example, if the system receives a request from two different clients to update a given record (i.e., conflict), the system, applying internal synchronization logic, can eliminate propagating the first update, as it is rendered moot by the second update. In this manner, the system can apply a first-level resolution of requests that are conflicting or complementary and, as a result, eliminate those synchronization activities which are redundant or moot.

B. Overview of synchronization system internal architecture

[0039] Fig. 2 is a block diagram illustrating a modular or high-level view of the synchronization system 200. As shown, the synchronization system 200 includes a synchronization engine (core) 230 that is connected to both a Grand Unification Database(s) (GUD(s)) 210 and to an action queue 240. As also shown, the engine presents two interfaces, a client API 220 and type API 250, for communicating with components outside the core engine.

[0040] The GUD 210, as previously described, serves as a central repository storing record data and mappings which dictate how records are transformed (i.e., from one data set to

another). The synchronization engine 230 includes generic logic for managing the GUD 210, including locating and interpreting information in the GUD. Based on the information in the GUD 210 and client requests, the synchronization engine 230 builds the action queue 240, adding or removing specific tasks from the queue as necessary for carrying out synchronization transactions. The action queue 240 itself is an array of task entries; it may grow or shrink, depending on the current number of entries that it stores. In the currently-preferred embodiment, the array is sorted by record ID, that is, according to the record ID of the corresponding record from the GUD. Since entries are sorted by record ID, the task of identifying entries in conflict is simplified.

10 [0041] To communicate with the clients, the synchronization engine 230 employs the client API 220. The client API provides database engine-like functionality. For example, API function calls are provided for moving to records, reading records, and writing records. In the currently-preferred embodiment, clients accessors 221, 223 are "accessor" portions of the synchronization system which, in turn, communicate directly with the "real" clients, such as
15 REX. By implementing its architecture such that all clients communicate commonly through the client API 220, the system 200 provides plug-in capability for supporting new clients.

[0042] In order for the system to correctly determine record information in the GUD 210, the synchronization engine 230 communicates with type drivers or modules (e.g., X type 251 and Y type 253) through the type API 250. As previously described, each type, such as calendar,
20 contacts, and the like, is associated with a particular type module. The type API 250 allows the synchronization engine 230 to ask common questions about information stored in the GUD 210. For example, if the synchronization engine 230 needs to determine whether two records are identical, it can request a record comparison operation by the corresponding type module, using the type API 250. In comparison to the client API 220, the type API 250 is comparatively small.

25 By implementing its architecture such that all type-specific requests are communicated commonly through the type API 250, the system 200 provides built-in extensibility. When support is desired for a new type, one need only plug in a new type module. Any client which wants to communicate with that new type now has automatically gained support for that new type. In the currently-preferred embodiment, a type module is unaware of any specific clients

which it supports. Clients, on the other hand, typically know what types that each desires to synchronize with.

[0043] As also shown, each client accessor can communicate directly with the type modules, using a record API 260. In the currently-preferred embodiment, each type module surfaces its own record API, such as record API 260 for type module 251. The underlying record API is specific for each type. Each accessor communicates with a desired type module, not through the synchronization engine 230, but instead through the exposed record API for the desired type. Thus, in effect, there is a direct communication path between client accessors and type modules. In typical use, the record API is employed by a client accessor to create or write record-specific information. For example, if the client desires to write a "subject" for a contact record, the client, operating through the corresponding client accessor, can invoke the corresponding record API for requesting this service. In response to invocation of the record API, the corresponding type module would service the API call for assisting with creating or editing the underlying record, in the matter requested by the client. The actual work of creating or editing the record is typically performed by the client; however, the corresponding type module returns specific information about the given type, so that the client knows exactly how the record is structured. As a simple example, the record API might return information indicating that a particular record type consists of a structure having four string data members, each being 64 bytes long. Based on such information, the client now knows how to interpret and process that type.

C. Synchronization system detailed internal architecture

1. GUD

[0044] Fig. 3 is a block diagram illustrating organization of a GUD 300. In the currently-preferred embodiment, the system implements one GUD per type. For instance, if one were synchronizing contacts, calendars, and "to do"s (i.e., task-oriented information), one would have three GUDs, one for each type. As shown, each GUD database internally stores two sets of tables: mapping tables 320 and data table 310. The data table 310 stores the actual record data 313 (i.e., blob data), together with a unique reference (ref) ID or "GUD ID" 311. In the presently-preferred embodiment, each reference ID (e.g., a 32-bit or 64-bit ID) is unique not only

within its particular GUD database but also across all GUD databases. Thus, for example, the system would not duplicate a calendar reference ID in the contact GUD database. With this approach, the individual data items are uniquely identified across the entire system. If desired, the GUD itself (or its data record portion) may be implemented as one of the actual client data sets (i.e., one of the data sets serves as the GUD, or portion thereof).

[0045] Also shown, mapping tables 320 store entries comprising a reference ID 321, a source ID 322, a checksum or integrity value (e.g., CRC) 323, and a last modification (mod) timestamp 324. The reference ID 321 is the same ID as associated with a record in the data table 310. The source ID 322 is the record ID for the record, as it was received from the client. The last modification timestamp 324 establishes when the record was last synchronized through the system. The timestamp (e.g., system time structure) reflects the time on the system clock of the machine which is being synchronized. Optionally, the system stores a comparison value or checksum (e.g., cyclic redundancy checking or CRC) 323, for use with those clients that do not support timestamps. If the checksum is not used, the system stores 0 as its value.

[0046] Each table itself is linked to a particular client, through a table ID, with the correspondence being stored as configuration information (which in the currently-preferred environment exists as a higher level than the synchronization engine). In this manner, each one of the mapping tables can be associated with an appropriate client. The end result is that the system maintains a mapping table for each client. Thus, for a given record ID, the system can easily determine (from the above-described reference ID-to-source ID correspondence) where that record maps to for all clients. Consider, for instance, a particular record residing on a REX device. Based on the source ID for that record, the system can determine from the mapping table the corresponding mapping table item for that source ID. Now, the system has sufficient information allowing the particular record to be synchronized, as required by the user. When the data is completely synchronized with all clients, all mapping tables in the system will store that record ID (i.e., the record ID is now common to all tables once the data is completely synchronized with all clients).

2. Action queue

[0047] The action queue stores entries of a particular action type, which are used during synchronization to indicate all actions needed to be performed by the system. In the currently-preferred embodiment, six action types are defined:

5

- (1) GUD_UPDATE
- (2) GUD_ADD
- (3) GUD_DELETE
- (4) CLIENT_UPDATE
- 10 (5) CLIENT_ADD
- (6) CLIENT_DELETE

10

The first three action types or "GUD action types" indicate actions to be performed against the GUD. For example, if the system receives a new record from a client, it must add the new record
15 to the (corresponding) GUD; this is indicated by an action queue entry having a type of GUD_ADD. In operation, the system will not only add the record to the corresponding GUD but, also, will eventually add that record to other clients which are associated with that record as well (unless the user instructs otherwise). In a similar manner, a GUD_UPDATE action item or command will result in the system updating the corresponding GUD for a given record (e.g., as a
20 result of that record having been modified at the client), and a GUD_DELETE action item or command will result in the system deleting the record from the corresponding GUD (e.g., as a result of that record having been deleted at the client).

20

[0048] The CLIENT action types are used to indicate particular synchronization work which is required to be performed for a particular client. Suppose, for instance, that the
25 synchronization engine determines that the REX client needs to be updated, as a result of actions undertaken by other clients; the REX client need not be currently available (e.g., need not be currently connected to the system). In such a case, the engine can post to the action queue appropriate action entries for indicating the synchronization work which is required to be performed the next time the REX client is connected. In a manner similar to that described
30 above for the GUD, the system can specify an update (CLIENT_UPDATE), add

25

30

(CLIENT_ADD), and/or delete (CLIENT_DELETE) action, on a per client basis. In the instance of an update or delete action, there already exists a corresponding mapping table item. For an add action, however, the system undertakes as its first action item the task of creating a new mapping table item. Therefore, when the add action is eventually performed, the table item will be created as well. On the other hand, should the action be canceled, the mapping table item will not be created.

[0049] Additional pieces of information are tracked by each entry in the action queue:

(1) record data, (2) source client, and (3) timestamp. The record data is the actual data (or a reference to the actual data) obtained from the client. In this manner, the actual data may be associated with a particular action. The source client indicates which client the action originated from. This is useful, for instance, during synchronization, so that the system does not attempt to synchronize the client from which the data just arrived. The timestamp stored in an action queue entry is the last modification time of the record from the source client. This is stored for possible use during conflict resolution (which is described in further detail below).

[0050] As previously described, the entries in the action queue are sorted by reference ID. In this manner, the system can quickly determine action queue entries which are potentially in conflict. For example, if the queue contains three entries all having the same reference ID, the system must examine those entries for uncovering any conflicts. The actual conflict resolution rules applied in the system are described below.

3. Methodology of system operation

[0051] Fig. 4A illustrates an overall methodology 400 of the present invention for providing synchronization contemporaneously among an arbitrary number of clients. At step 401, the system initializes all clients and types (data structures). At step 402, the system establishes a loop for determining for each client what actions are to be performed. Here, the system begins building the action queue. Once the action queue or table has been built, the system proceeds to resolve any conflicts present. This is indicated by step 403. In particular at this step, the system performs housekeeping on the queue, removing any action entries which are unnecessary.

[0052] Conflict resolution requires further explanation. As previously described, the entries in the action queue are sorted by reference ID. In this manner, the system can quickly determine action queue entries which are potentially in conflict. For example, if the queue contains three entries all having the same reference ID, the system must examine those entries for uncovering any conflicts. Not only are items in the action queue sorted by a reference ID but, as a second level of ordering, they are also sorted by action. GUD updates are always sorted to the top, thus establishing their priority over other types. Now, the following exemplary conflict resolution rules may be applied:

10 Rule 0:

+ GUD_UPDATE
+ <entry(ies) other than GUD_UPDATE>

GUD_UPDATE wins; delete all others

15

Rule 1:

+ GUD_UPDATE
+ GUD_UPDATE

GUD_UPDATE with greatest timestamp wins (or display UI)

20

Rule 2:

+ GUD_UPDATE
+ GUD_DELETE

GUD_UPDATE (take data over non-data)

25

Rule 3:

+ CLIENT_UPDATE
+ CLIENT_UPDATE (from another client)

Leave both (i.e., same)

30

Once conflicts have been resolved the action queue is ready for use. Specifically, at step 404, the system processes all remaining action entries in the action queue. The actions themselves are performed on a transaction-level basis, where a transaction comprises all actions performed on a given record GUD ID. Thereafter, the system may perform cleanup, including closing any open databases and freeing any initialized data structures (e.g., type).

35

[0053] Fig. 4B illustrates particular substeps which are performed in conjunction with step 402. The substeps are as follows. At step 421, the system determines all updates and adds originating from the client (i.e., the client currently being processed during the "for" loop). In essence, the system operates by asking the client for all modifications (e.g., updated or added records) since last synchronization. Once these are learned, the system places them in the action queue, either as a GUD_UPDATE or GUD_ADD. If desired, a filter may be applied at this point, for filtering out any records which are desired to be omitted from the synchronization process. The next step, at step 422, is for the system to determine any deletions coming from the client. Note, here, that the update/add step (421) comes before the deletion determination step (422). This allows the system to determine what is new before determining what has been deleted. As an optimization at this point, the system can look at the record count at the client for determining whether in fact there have been any deletions at all. In the event that the count indicates no deletions, the system can eliminate the time-consuming process of determining deletions (which may require the system to examine numerous records individually). At step 423, the system makes a reverse determination: determining any updates or adds which need to be sent from the GUD back to the client. The mapping table stores a timestamp indicating when the client was last synchronized as well as a timestamp for each record item. Accordingly, the system can determine whether the item needs to be updated or added at the client. In the currently-preferred embodiment, the timestamp is generated based on the system clock of the client which is undergoing synchronization. Finally, at step 424, the system determines any deleted records in the GUD, for indicating which corresponding records should be deleted from the client. Specifically in the mapping table, each entry includes a deletion flag which may be set for indicating deletion of the corresponding record. These foregoing steps are performed for all clients undergoing synchronization, until the action queue is filled with the appropriate action entries required for effecting synchronization.

[0054] Fig. 4C illustrates particular substeps which are performed in conjunction with step 404. The substeps are as follows. At step 431, the system determines whether the action is from one client to another client. If the action is to a client, the system may simply proceed to update the client, as indicated by step 432. If, on the other hand, the action is from a client, the system must update the GUD, as indicated at step 433, and, in turn, propagate the update to the

other clients, as indicated at step 434. The actual propagation is performed recursively invoking itself as client actions (rather than GUD actions). Here, the system fabricates a surrogate or fake action item which is then acted upon as if it were from the action queue. All the time during the method, the GUD has played an important role as a data source for those clients which are not
5 currently available.

[0055] While the invention is described in some detail with specific reference to a single-preferred embodiment and certain alternatives, there is no intent to limit the invention to that particular embodiment or those specific alternatives.

WHAT IS CLAIMED IS:

5 1. (currently amended) In a data processing environment, a method for synchronizing multiple data sets, the method comprising:

 establishing a data repository for facilitating synchronization of user information maintained among more than two data sets, said data repository storing user information that is a super-set of all user information for which any user desires synchronization support;

10 storing at least one mapping which specifies how user information may be transformed for storage at a given data set;

 receiving a request for synchronizing at least one data set;

 based on user information stored at said at least one data set and based on said at least one mapping, propagating to the data repository from ~~[[each-of-at]]~~ said at least one data set any changes made to the user information, to the extent that such changes can be reconciled with user information already present at said data repository; and

15 based on user information stored at said data repository and based on said at least one mapping, propagating to ~~[[each-of]]~~ said at least one data set any changes to the user information which have been propagated to the data repository, to the extent that such changes are not present at said ~~[[each]]~~ at least one data set.

20 2. (original) The method of claim 1, wherein said step of propagating to the data repository comprises:

 performing selected operations of adding, updating, and deleting information at the data repository, so that the data repository reflects changes made to user information at the data sets.

3. (original) The method of claim 2, wherein said operation of deleting information comprises a logical delete operation of marking information as having been deleted.

4. (canceled)

5. (original) The method of claim 1, wherein said data repository and said at least one mapping comprise a grand unification database, for facilitating synchronization among multiple data sets.

6. (original) The method of claim 5, wherein one grand unification database is created for each type of user information which is to be synchronized.

7. (original) The method of claim 6, wherein said environment includes types of user information selected from contact, calendar, and task-oriented information.

8. (canceled)

9. (original) The method of claim 1, wherein each data set comprises a plurality of data records, and wherein each data record is represented within the data repository.

10. (original) The method of claim 9, wherein each of said data records is represented within the data repository by a corresponding data record having a unique identifier.

11. (original) The method of claim 1, wherein each mapping comprises a mapping table storing a plurality of mapping entries, each mapping entry storing at least a first identifier for indicating a particular data record in the data repository which the entry is associated with, and a second identifier for indicating a particular data record at a particular data set which is the source for the user information.

12. (original) The method of claim 11, wherein each mapping table is associated with a particular data set.

13. (original) The method of claim 11, wherein each mapping entry stores particular information useful for determining when its associated user information was last modified.

14. (original) The method of claim 13, wherein said particular information comprises a last-modified time stamp, derived at least in part from the client device where the associated user information was last modified.

15. (original) The method of claim 13, wherein said particular information comprises a checksum value, for use with a data set residing at a client device that does not support time stamps.

16. (original) The method of claim 1, wherein said step of propagating to each of said at least one data set comprises:

performing selected operations of adding, updating, and deleting information at each of said at least one data set, so that said each reflects changes made to user information at other data sets.

17. (original) The method of claim 16, wherein said operation of deleting information comprises physically deleting information at said each data set.

18. (original) The method of claim 1, wherein at least one of the said data sets functions, at least in part, as said data repository.

19. (original) The method of claim 1, wherein user information is stored at the data repository as unformatted blob data.

20. (original) The method of claim 19, further comprising:

providing at least one type module for facilitating interpretation of user information stored as unformatted blob data at the data repository.

21. - 40. (canceled).

DATA PROCESSING ENVIRONMENT WITH METHODS PROVIDING
CONTEMPORANEOUS SYNCHRONIZATION OF TWO OR MORE CLIENTS

5

ABSTRACT OF THE DISCLOSURE

A synchronization system providing multi-client synchronization is described. By storing the data that is actually being synchronized (i.e., storing the actual physical body of a memo, for instance) inside an extra database, "Grand Unification Database" (GUD), (or by
10 specially-designated client data set) under control of a central or core synchronization engine, rather than transferring such data on a point-to-point basis, the system of the present invention provides a repository of information that is available at all times and does not require that any other synchronization client (e.g., PIM client or hand-held device) be connected. The GUD provides a super-set of the other client data sets. Therefore, if the user now includes an
15 additional client, such as a server computer storing user information, the synchronization system has all the information necessary for synchronizing the new client, regardless of whether any of the other clients are currently available. The system can, therefore, correctly propagate information to any appropriate client without having to "go back" to (i.e., connect to) the original client from which that data originated.